

A Security Mediator for Health Care Information

Gio Wiederhold, Ph.D., Michel Bilello, Ph.D.
Stanford University, CA

Vatsala Sarathy, M.S., XiaoLei Qian, Ph.D.
Oracle Corp., Redwood City, CA, SRI International, Menlo Park, CA

July 25, 1996

The TIHI (Trusted Interoperation of Healthcare Information) project addresses a security issue that arises when some information is being shared among collaborating enterprises, although not all enterprise information is sharable. It assumes that protection exists to prevent intrusion by adversaries through secure transmission and firewalls. The TIHI system design provides a gateway, owned by the enterprise security officer, to mediate queries and responses. The latter are typically transmitted via the Internet. The enterprise policy is determined by rules provided to the mediator. We show examples of typical rules. The problem and our solution, although developed in a healthcare context, is equally valid among collaborating enterprises.

INTRODUCTION

We address an issue in the protection of medical information that is starting to arise as the basic infrastructure for secure transmission and storage enters into practice. We assume an environment where encrypted transmission, firewalls, passwords, and private and public keys provide adequate protection from adversaries. The problem which remains, and addressed here, is now to enable selective sharing of information with collaborators, without the risk of exposing related information in one's enterprise domain or enclave that needs to be protected [1]. It is assumed that remote information sharing is carried out via the Internet. We will first sketch some examples to clarify the problem and then formulate the model for our work.

In a hospital the medical record system collects a wide variety of information on its patients. Most information on a patient must be accessible to the clinical healthcare personnel, including community physicians, and a substantial fraction to the hospital billing clerks [2]. Similar data are requested

by insurance companies, and certain data and summarization are necessary for hospital accreditation and public health monitoring. Results for all of these customers must be handled distinctly.

In a manufacturing company collaborations are often formed with suppliers and marketing organizations. Such virtual enterprises are formed to design, assemble, and market some specific products. Design specifications and market intelligence must be rapidly shared to remain competitive. These collaborations overlap, producing security problems which are stated to be the primary barrier to the acceptance of this approach [3]. Uncontrolled sharing of proprietary data is too risky for a manufacturer to grant a supplier. The supplier will also be wary of giving information to the customers.

These two scenarios have the following commonality.

1. We are dealing with friends, not enemies, and should provide relevant information expeditiously.
2. The collected information is not organized according to the needs of a security protocol.
3. It is impossible to rigorously classify the data, a priori, by potential recipient.
4. It cannot be fully determined from the query whether the results combine information which should be withheld.

For instance, a medical record on a cardiac patient can include notations that would reveal a diagnosis of HIV, which should not be widely revealed, and withheld from cardiology researchers if HIV status is irrelevant to their research. A design document on a plastic component, to be outsourced, also indicates the incorporation of a

novel component supplied by another manufacturer, which provides a competitive advantage.

Our model formalizes the role of a security officer who has the responsibility and the authority to assure that no inappropriate information leaves an enterprise domain. A firewall protects the domain vis-a-vis invaders. Distinct gateways, each owned and controlled by a security officer, provide the only legitimate pathways out of, and into, the domain. This gateway is best envisaged as a distinct computer system; we refer to such a system as a security mediator, placed as sketched in Figure 1. In the security mediator the policies set by the enterprise on security and privacy are implemented, under control of, and through interaction with the security officer. Databases and files within the domain can provide services and meta-data to help the activities of the security mediator, but cannot be fully trusted. The security mediator is able to use secure communication and authentication of outside requests.




Figure 1: Security mediator setting.

It is important to recognize, as sketched in Figure 2, that validation of communication content must occur both with respect to the query and the responses. For instance, it is inadequate to allow a validated researcher in cardiac diseases to receive all records on cardiac patients, if that also includes HIV cases. Depending on institutional policy, such cases will be omitted or sanitized.

ARCHITECTURE

The mediator system consists of modules that performs the following tasks:

- Processing of query (*pre-processing*)
- Communication with databases (submission of query and retrieval of results)
- Processing of results (*post-processing*)




Figure 2: Two-way information content validation.

- Writing into a log file

The system is best visualized as residing on a distinct workstation, operated by the security officer. Within the workstation is a rule-base system which investigates queries coming in and responses to be transmitted out. Any query and any response which cannot be vetted by the rule system is displayed to the security officer, for manual handling. The security officer decides to approve, edit, or reject the information. An associated logging subsystem provides both an audit trail for all information that enters or leaves the domain, and provides input to the security officer to aid in evolving the rule set, and increasing the effectiveness of the system.

The mediator system can operate fully interactively or partially automatically. A reasonable goal is the automatic processing of say, 90% of queries and 95% responses, but even an empty, rule-less security mediator can greatly improve operations. Without rules, all interactions are presented to the security officer, but they are now viewed online, and immediately editable, if needed. In time, simple rules can be entered to reduce the load on the officer. Currently, paper-based systems often take weeks for turn-around, or are bypassed. Even when operating automatically, the security mediator remains under the security officer's control. For example, rules are modifiable by the security officer at all times. In addition, daily logs are accessible to the officer, who can then keep track of the transactions.

The rules balance the need for preserving data privacy and for making data available. Data which is too tightly controlled would be less available and useful for outside users. Conversely, a sufficient level of protection of data privacy must be main-

tained.

The mediator system and the medical database typically reside on different machines. Thus, since all queries are processed by the mediator, the database need not be secure unless it operates in a particularly high security setting.

THE RULE SYSTEM

In order to automate the process of controlling access and ensuring the security of information, the security officer enters rules into the system. The security mediator uses these rules to determine the validity of every query and make valuable decisions pertaining to the dissemination of information. The system helps the security officer enter appropriate rules and update them as the security needs of the organization change.

The rules are simple, short and comprehensive. They are stored in a database (distinct from the main medical database) with all edit rights restricted to the security officer. Some rules may be related to others, in which case the most restrictive rule automatically applies. The rules may pertain to users, sessions, tables or any combinations of these.

Once they are entered into the system by the officer, all the rules will be checked for every query issued by the user in every session. All applicable rules will be enforced for every user and the query will be forwarded only if it passes all tests. Unless a rule permits explicit pass through, it goes to the security officer. In the event a rule is violated by a query, the error message will be directed to the security officer and not to the end user. The users in such cases will not see the error message. This is necessary because even error messages could be interpreted and meaningful inferences could be made, or the user could rephrase the query to bypass the error. The errors as well as all queries will be logged by the system for audit purposes.

The security mediator checks outgoing results as well. This is crucial since queries are inclusive, not exclusive selectors of content and may retrieve unexpected information. Thus, even when the query has been validated, the results are also subject to screening by a set of rules. As before, all rules are enforced for every user and the results are accessible only if they pass all tests. Also, if the results violate a rule, an error message is sent to the security officer but not to the user.

Not only are the rules easy to comprehend and to enter into the system, they are also powerful enough to enable the officer to specify requirements

and criteria accurately, so that whenever users may see all information, they should be allowed to do so and whenever information is restricted, they should have no access to it. The users in the system are grouped as cliques and rules may apply to one or more cliques. The security officer has the authority to add or delete users from cliques and to create/drop cliques. Similarly, columns in tables can be grouped into segments and query/results validations could be performed on segments.

The rules can be classified as *set-up* or maintenance rules, *pre-processing* (query) rules and *post-processing* (result) rules. Some rules may be both *pre-* and *post-processing* rules. Examples of *pre-processing* rules include limitation of the number of queries per session, the number of queries for the clique, the session duration, etc. *Post-processing* rules include a minimum number of rows retrieved, restrictions on intersection of queries. A more comprehensive list of rules can be found in the appendix.

SYSTEM OPERATION

The following sequence of operation is applied for every request.

- When the user solicits the system, the security mediator assesses clique membership. The user enters a query.
- Next, the mediator parses the query. If parsing is not successful, an error message is sent out to the security officer.
- Then, it checks to see if access to all columns specified in the SELECT and WHERE clauses in any segment is permitted to the members of the clique. If not, an error message is sent to the security officer.
- It then looks at every *pre-processing* rule in the system and validates the query against each. If any rule is violated, an error message is sent to the security officer.

At this point, the query is actually processed and results are obtained by the mediator.

- Now the *post-processing* rules are applied.
- On textual results, rules may specify that all words must come from a specified vocabulary. Any unknown term will be presented, with surrounding context, to the security officer, and if not approved, no result will be returned.

- Lastly, further result modification is done as specified by the rules. Operations that can be invoked include random falsification of data and aggregation.
- Security officers can edit documents brought to their attention before releasing them. That should include whiteing-out portions of graphics and design drawings.
- Now the results are sent back to the user. Then the mediator updates internal statistics such as number of queries for the session, duration of the session, etc. It also updates the log files appropriately. This last step is done in all cases, whether or not there were errors.

VIEW-BASED ACCESS CONTROL

Notice that the tables referred to in the rules do not have to be base relations. They can be derived relations or views defined by arbitrary Structured Query Language (SQL) queries. Hence, the set of rules collectively specifies a view-based access control policy.

Views in relational databases have long been considered ideal as the objects of access control, because they have a higher degree of logical abstraction than physical data and hence enable content-based or context-based security, as opposed to container-based security provided in operating systems.

View-based access control in relational databases was first introduced in IBM's System R [4], in which views expressed in SQL are the objects of authorization. It has been adopted by most commercial relational DBMSs. However, view-based mandatory access control has not been in widespread use because of the security problem [5]. The security question asks the following. Is there a database state in which a particular user possesses a particular privilege for data in a specific view? In container-based access control, different containers do not share contents. Hence, a secret label on a container guarantees that data in the container are not accessible to unclassified users. In view-based access control, however, views might overlap because the same data might satisfy more than one view. Hence, a secret label on a view does not guarantee that data contained in the view are not accessible to unclassified users.

To support view-based mandatory access control, queries have to be analyzed and answers have to be filtered to ensure that data in a view are accessible to all and only to those users who are au-

thorized to access the view. We envision two types of query analysis.

It turns out that the security mediator should not accept a query solely on the condition that the issuer has authorization for all relations mentioned in the query, base or derived. Instead, the security mediator should try to reformulate the query using those views for which the issuer of the query has authorization. If a reformulation is possible, then the reformulated query will be evaluated in place of the original query. Otherwise, the original query is rejected. This approach will also facilitate the evolution of the security policy enforced by the security mediator.

Note that access control on a per-query basis might not be sufficient. Even when a user has authorization to every query issued, he might be able to combine answers from a sequence of queries to derive data in a view for which he does not have access authorization. Such scenarios necessitate the need for the security mediator to keep track of the access history for every clique/user.

CONCLUSION

We are addressing privacy and security maintenance in collaborative settings, where information has to be selectively protected from colleagues, rather than withheld from enemies. The problem only arises once a basic secure infrastructure is established. Today, privacy protection in healthcare is preached, but ignored in practice, putting many institutions at risk. In crucial settings, security officers control input and output, but do so on paper, so that interactions are typically delayed by weeks, and high costs are incurred due to delays and misunderstandings. The primary barrier, as stated by Hardwick, to the realization of virtual enterprises is insufficient security controls. The corporations participating in a virtual enterprise are independent and frequently compete against one another.

The approach we are developing provides tools for a security officer. Database systems have provided tools to control queries, under the aegis of the database administrator. We illustrated above that query-only tools are inadequate in complex settings, and we emphasized the need for view-based access control. In addition, the major role of a database administrator is to help customers get maximal relevant data, a task that often conflicts with security concerns. Furthermore, the majority of data is not in database systems that provide security, and even less data resides in costly, validated multi-level secure systems.

The concept of security mediator as an intelligent gateway protecting a well-defined domain is clear, simple, and the cost of modern workstations make it feasible to assign such a tool to a security officer. Like most security measures, the security mediator cannot offer a 100% guarantee, especially with respect to statistical data security. But having a focused node, with a complete log of requests and responses, and an incrementally improving rule collection, provides a means to ratchet protection to a level that serves the healthcare institution or other enterprise needs and policies effectively.

APPENDIX

Examples of set-up rules

set logfile "x"

The table or path name to the log file

create clique x

Create a clique of users called x

add user user_name clique_name

Add user called user_name to clique_name

drop clique x

create segment segment_name

Examples of pre-processing rules

set clique stat_only true/false

Only statistical info (average, median)
allowed for user

limit queries_per_session x

Number of queries allowed in a session

limit clique queries x

For a given user, number of queries allowed
per session

limit clique segment

Limit all users in clique to columns/tables in
segment. This specifies explicit pass through
of results.

Examples of post-processing rules

set user table random on/off

Random falsification of data to be
performed or not for user/table combination

limit min_rows_retrieved x

Minimum number of matching rows for a
given selection criterion

secure keyword segment_name

Do not pass queries whose results contain
the keyword in the segment.

limit clique intersection x

No two queries by user can have an
intersection greater than x rows

Acknowledgments

This work was supported in part by NSF grant ECS-94-22688.

References

1. Johnson DR, Sayjdari FF, Van Tassel JP. Missi security policy: A formal approach. Technical Report R2SPO-TR001, National Security Agency Central Service, July 1995.
2. Braithwaite B. National health information privacy bill generates heat at SCAMC. *Journal of the American Informatics Association*, 1996:3(1):95-96.
3. Hardwick M, Spooner DL, Rando T, Morris KC. Sharing manufacturing information in virtual enterprises. *Comm. ACM*, 1996:39(2):46-54.
4. Griffiths PP, Wade, BW. An authorization mechanism for a relational database system. *ACM Transactions on Database Systems*, 1976:1(3):242-255.
5. Schaefer M, Smith G. Assured discretionary access control for trusted RDBMS. In *Proceedings of the Ninth IFIP WG 11.3 Working Conference on Database Security*, 1995:275-289.